ROBERT FRANCES GROUP
Business Advisors to IT Executives

Apr. 21, 2023

IBM Turns Up the Volume on Hybrid Cloud and AI Capabilities for IBM zSystems and LinuxONE

By Jean S. Bozman

The rubber has finally met the (distributed) road.

For datacenters where the IBM zSystems are playing a leading role, IBM is evolving its mainframe platform to enhance hybrid cloud and AI/ML workloads. As it does that, IBM is giving customers the option of installing new IBM z16 systems and LinuxONE systems in single-frame enclosures made by IBM -- or in industry-standard rack-mounted enclosures that at IBM customers' sites.

This deployment model, offering more options for workload placement, is designed to increase on-board hybrid-cloud and AI capabilities, while reducing data-center power and energy requirements for the virtualized workloads that run customers' core business.

The move is designed to reduce power and cooling requirements in data centers and in computing hubs, supporting customers' power and cooling and systems management goals. With its April announcement, IBM said its LinuxONE Rockhopper systems could reduce energy consumption by 75%, and space by nearly 70%, compared with other deployments of large numbers of x86 server platforms running similar Linux applications.

For IBM, this consolidation model for gathering Linux workloads on LinuxONE and IBM zSystems platforms has become a stated point-of-differentiation regarding the way IBM systems run, manage, and secure Linux workloads.

In doing so, IBM is showing how it is delivering on its core strategy of focusing on both hybrid cloud (HC) and AI as two leading paths for customers who are transforming their infrastructure to meet a wide variety of Core, Cloud and Edge imperatives.

Co-location deployments give customers the option to manage hybrid cloud and Red Hat distributed containers alongside highly scalable enterprise

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; Contact: inquiry@rfgonline.com

computing. With the z16's on-board AI acceleration, businesses could speed up business analytics for faster actionable insights based on corporate data.

Evolving the Operational Model for Linux Enterprise Computing

As the digital economy transforms many businesses, IBM supports the transactional compute models with its IBM z16, introduced last April, as it moves to support an expanding base of customers who deploy container-based applications across their business units, enterprise-wide.

With this announcement, IBM is offering customers the option of installing IBM zSystems in IBM-manufactured single-frame enclosures – or in customers' onsite rack enclosures.

Specifically, the z16 single-frame and the LinuxONE Rockhopper 4 are available for installation in two formats:

- The z16 A02 can be shipped in an IBM rack and the LinuxONE Rockhopper 4 can be shipped in an IBM rack.
- The z16 AGL can be installed in a non-IBM customer rack and the LinuxONE Rockhopper can be installed in a non-IBM customer rack.

Either approach supports consolidation of Linux workloads, because thousands of Linux applications can run inside an IBM z16 platform – or inside one or more IBM LinuxONE platforms. The difference now is that customers will have the choice of installing these systems either in single-frame enclosures made by IBM – or in customers' standard rack enclosures that were made by other vendors.

Both types of enclosures allow customers to consolidate Linux applications that may have been running on large numbers of x86 servers before.

Importantly, the newly announced installation modes for the IBM z16 and LinuxONE Rockhopper 4 systems may open the door for independent VARs and MSPs to take on that responsibility in the future. Initially, IBM personnel will install the first generation of the new LinuxONE systems in non-IBM racks at customer sites. However, we believe it is entirely possible that non-IBM personnel will eventually take on that task, following a certification course.

IBM Strategy for HC and AI

All of this is consistent with IBM CEO Arvind Krishna's 2020 speech when he took office that year – and in the years since then – that HC and AI would be two primary pillars of IBM's strategy for years to come.

Prior to becoming CEO, Krishna was IBM Senior Vice President (SVP) for Cloud and Cognitive Software, where he drove the company's acquisition of Red Hat. It is Red Hat Open Shift that enables much of the consolidation of virtualized workloads – and ongoing management of those workloads via Ansible.

In January, 2023, Krishna called hybrid cloud and AI "the two most transformative technologies for business today," as he said that IBM is well positioned to take advantage of customers' continued investments in both of these technology areas.

Now, IBM is marketing the IBM z16 Systems and LinuxONE platforms together to industry segments that require high levels of availability, security, and predictability. They are working to blend two computing styles: frame-based installations for data-center deployments, and rack-based installations designed to gather together, and consolidate, highly distributed Linux workloads for operational savings in power, cooling and floor-space.

Called out by name in the announcement are these industry segments: financial services, healthcare, and transportation – all of which have extensive

government compliance requirements, across multiple geographies, worldwide. IBM believes it has a competitive advantage in this space, by providing end-to-end pervasive encryption on IBM zSystems – and extending it to Cloud and Edge locations using IBM software to enforce end-to-end confidential computing.

This follows onto the April, 2022, announcement that the z16 mainframe, based on the IBM-designed 7nm Telum processor, supported AI inferencing at-scale, and quantum-safe cryptography. IBM announced its LinuxONE Emperor 4 system in September, 2022, highlighting its scale, performance, and security, along with power/cooling and floor-space improvements to address sustainability and ESG requirements.

Analyst Perspective

During the three years since Arvind Krishna became IBM CEO, IBM executives in IBM's Systems, Cloud and Research groups have been describing these specific industry segments – finance/banking/insurance, healthcare, and transportation – as benefiting from RAS and end-to-end security features across the IBM z product lines.

While IBM is not alone in promoting end-to-end performance, reduced latency, and faster processing, it is identifying the z and LinuxONE platforms as key to enabling its HC + AI strategy. But IBM has the advantage of guaranteeing enterprise servers that provide 99.99999 percent availability (seven-nines of availability, with less than five seconds of outage per year).

For distributed computing, this approach will protect data and accelerate business analytics from the Core to the Edge, leveraging IBM zSystems and LinuxONE strengths in RAS (reliability, availability, and serviceability) to benefit containerized applications and distributed data.

Clearly, other industry segments outside the financial/banking sector will leverage these IBM zSystems and LinuxONE features to improve hybrid-cloud and multi-cloud performance, while managing end-to-end enterprise computing requirements. Examples include pharmaceutical companies, manufacturing companies, and energy/oil/gas companies – all of which currently use containerized, distributed applications to run mission-critical business units.

The Enterprise Role of Linux Systems

It is not surprising that Linux has a prominent role in this IBM strategy.

Linux, after all, has become the *lingua franca* of the Internet and cloud computing, while enterprise applications and data must not only be *open* but also *secured.* Examples of industries that will adopt secured Linux – across the enterprise from Data Center to Cloud to Edge – include pharmaceutical companies, manufacturing companies, and energy/oil/gas companies.

It's worth noting that scalable IBM z16 systems, using their Integrated Facility for Linux (IFLs) for application isolation, can support thousands of Linux developers working simultaneously. This has been known for many years, and was initially demonstrated 20 years ago at IBM's Boeblingen, Germany, location. LinuxONE Rockhopper, whether installed in the customer data center or in regional locations, is capable of consolidating thousands of Linux instances – and thousands of workloads.

In the datacenter, reducing costs relies on reducing server footprints, as was proven in the early 2000s, as virtualization significantly increased the number of VMs per server. With this announcement, IBM is pointing out the business value of workload consolidation per "footprint" when it comes to power/energy savings and ESG considerations in customers' data centers.

Summary

Taken together, IBM z systems and LinuxONE systems carry forward the corporate strategy of hybrid cloud + AI acceleration for enterprise-wide applications and data. Red Hat OpenShift for containers, IBM data-management software, IBM storage software, and IBM Consulting support for AI will all play active roles in these enterprise-wide customer deployments.

Customers clearly want to have more consistent deployments and management of mission-critical and business -critical systems, blending cloud-style computing (including public clouds and private clouds) with transactional applications that grew up in the data center. This will apply to what some have called “sovereign clouds” that operate within countries to confirm with country-specific privacy, security, and ESG/sustainability regulations.

While most companies continue to have a multi-vendor, multi-cloud IT environment – and will so do for years to come – IBM has found a way to implement its core technical strengths, embedded in its z16 and LinuxONE platforms, and to apply them across the enterprise.

---

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; Contact: inquiry@rfgonline.com

46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; Contact: inquiry@rfgonline.com